

Highly Autonomous Systems Workshop

Richard Doyle, Robert Rasmussen, Guy Man and Keyur Patel
Jet Propulsion Laboratory

Overview

In our lifetime, through the eyes of simple robots, grand vistas on other worlds have been unveiled for the first time. Enigmatic questions compel us to go further, to touch these distant landscapes and learn the secrets of the solar system. Yet in trying, we find our reach wanting, limited by the link to Earth upon which our probes depend. We are learning that to explore further, these probes must go alone. And to go alone, they must become much more intelligent.

The Highly Autonomous Systems Workshop, held in Pasadena, California on April 9-10, 1997, celebrated the birth, in both fact and fiction, of this new generation of explorers. Our goal was to bring together visionaries and skeptics, practitioners and researchers in artificial intelligence, planetary and space science, spacecraft design, mission design and mission operations, along with partners in aerospace to discuss the important advances in autonomous systems that are propelling this genesis. In exploration, this technology encounters a true and natural test of its maturity where there is no quarter for mediocrity, but where all can freely watch and weigh its performance.

Our workshop was one of three symposia in 1997 honoring the fictional birthyear of the HAL-9000 computer. The other two were held at MIT and at the University of Illinois in Urbana, which is also the fictional birthplace of HAL-9000. Our meeting was distinguished by reports on actual work towards building intelligent spacecraft.

Representing government were attendees from NASA, DoD and DARPA. Aerospace attendees included Lockheed-Martin, Boeing and TRW. Academia was represented by institutions such as MIT, CMU, Georgia Tech.

The meeting was organized around technical sessions on history and visions for spacecraft autonomy, state-of-the-art autonomous systems, autonomy technology for 2001, and long-term challenges and benefits of autonomy. Each session concluded with a panel comprised of all speakers of that session.

Featured speakers included one of the founders of AI, Prof. Marvin Minsky of MIT and Dr. Louis Friedman, Executive Director of the Planetary Society. The symposium banquet was graced by two exceptional speakers: the creator of

HAL-9000, Dr. Arthur C. Clarke, and Prof. David Stork of Stanford University, author of "HAL's Legacy."

The meeting concluded with the announcement of a university design competition on the intriguing topic of aerobots, an space platform design which operates in a planetary atmosphere, combining aspects of orbiting platforms and surface vehicles.

The workshop was hosted and organized by the authors and was sponsored by the NASA Autonomy and Information Management Program and the New Millennium Program Autonomy Integrated Product Development Team.

It is our aim by launching a series of workshops on the topic of highly autonomous systems to reach out to the larger community interested in technology development for remotely deployed systems, particularly those for exploration. We invite members of this community to join us to help guide and nurture autonomy's development, to learn about its enormous potential, both in space and here at home, to share ideas, and find a way to participate.

Spacecraft Autonomy: History and Visions

The first session took a historical view and described a set of visions for spacecraft autonomy that have arisen from different perspectives and evolved over many years of spaceflight.

Mark Brown of JPL described how deep-space missions have always had drivers for autonomy because of the impracticality of near-continuous communication and the unique difficulties associated with light-time delayed communication. Examples of long-standing drivers for autonomy on spacecraft include surviving failures, correct execution of time-critical activities (such as achieving orbit), onboard control requiring feedback, and protection of critical resources. A key concept is that when faults occur, the spacecraft must end up in a predictable, commandable state. Historically, autonomy has been applied only when deemed necessary, with onboard computing resources being a significant limiting factor.

Bob Connerton of NASA Goddard Space Flight Center spoke to the requirements for autonomy on spacecraft which observe the Earth from orbit. The overwhelming driver is the need to reduce up to terabits of raw data collected over a diverse set of high throughput space-based instruments to usable information, sometimes in near real-time. Onboard feature extraction and data fusion are important capabilities which can support responses to events such as volcanic eruptions or forest fires. Spacecraft will also be arrayed in formations and constellations. These multiple space elements must be precisely controlled and their onboard activities coordinated.

Dave Linick of JPL offered that, from an operations perspective, achieving autonomy involves migrating functions from the ground to the spacecraft. Indeed, a flexible flight / ground architecture would support functions being migratable in either direction, even during a single mission. Potential benefits include reduced costs, new kinds of missions that involve performing science investigations in uncertain environments, and relief on oversubscribed resources such as the Deep Space Network system of ground-based antennas for communicating with spacecraft. A clear driver is NASA's goal to have many more missions flying at once, via smaller, less expensive spacecraft, supported by a shared operations team on the ground.

Louis Friedman of The Planetary Society gave a talk entitled "Humans vs. Robots, or Where Will the Humans Be?" in which he argued for the value of unmanned spacecraft in performing the basic NASA mission of space exploration. Unmanned spacecraft will always be our first emissaries to remote places, and by making them more autonomous, they can perform new kinds of missions, more in-depth studies, and extend our scientific awareness further and further out. In this exciting picture of exploration, a theme is emerging which focuses on the search for life elsewhere: on early Mars, in the suspected subsurface ocean on Jupiter's moon Europa, at planets around nearby stars.

Matthew Barry of the United Space Alliance addressed autonomy from the perspective of NASA's manned spaceflight program. Risk management is the overriding consideration, with human lives being central to the picture. Nonetheless, with cost reduction becoming a major goal, there is considerable interest in capabilities like decision support, to assist flight controllers and/or astronauts in making procedural choices. In this context, fault diagnosis may be actually less important than sensible reconfiguration decisions, especially those requiring real-time or near real-time responses. Another area which would benefit greatly from automation support is mission planning and replanning.

Brian Williams of NASA Ames Research Center discussed a set of ideas for developing model-based autonomous systems. Recalling an episode from *Star Trek* in which Mr. Spock's brain is stolen for the purpose of becoming the control center for an entire planet, he noted that some current goals for using model-based reasoning to support onboard fault diagnosis and recovery were in fact similar in form if not scale. His Mode Identification and Reconfiguration (MIR) system continuously monitors qualitative representations of sensor data, identifying current spacecraft modes or states, and when these are fault modes, selects recovery actions based on a state transition diagram to return the spacecraft to the desired state.

Marvin Minsky of MIT gave a talk entitled "What Made HAL Late for His Party?" in which he lamented the lack of definitive progress on questions

which the field of Artificial Intelligence has long sought to address. He discussed theories of intelligence which address architectural issues, including his "Society of Mind." He suggested that a theory of which AI techniques actually work, and how well, in which domains, was probably achievable at this point. He also asserted that a key ingredient for success in NASA's efforts on spacecraft autonomy would be the development of comprehensive models and knowledge bases for onboard use, for both engineering and science purposes.

State-of-the-Art Autonomous Systems

The second session of the workshop surveyed current autonomous systems work where deployment has already taken place or is well-defined and imminent.

Perry McCarty of Lockheed-Martin reported on Autonomous Control Logic for autonomy on underwater submersibles. The basic challenge is to develop an onboard decision-making capacity which can continue a mission in the face of unanticipated, perhaps partially compromising events. A layered architecture combines a reactive component which monitors conditions and responds to anticipated events and a deliberative component which evaluates vehicle state and capabilities and chooses among courses of action with highest value towards completing the mission, even when the vehicle is found to be in a degraded condition.

Bruce Bullock of ISX and Dave Smith of Lockheed Martin spoke on the well-known Pilot's Associate Program. The Pilot's Associate is a real-time support system whose job is to efficiently enhance the situational awareness of a pilot in a tactical air battle situation. The task involves modeling the pilot's intentions and state of knowledge and inferring the intents and state of knowledge of friendlies and threats -- all while supporting the pilot's actions and communicating information and options accurately and unobtrusively. Plan generation and understanding, information management and real-time performance are key aspects of this complex human-machine system concept.

Doug Bernard of JPL reported on joint work between NASA Ames and JPL on the Remote Agent, which will be flight tested on the New Millennium DS1 mission in 1998. The Remote Agent Experiment will demonstrate an autonomy architecture consisting of three reasoning engines and associated knowledge bases: the Planner/Scheduler, which translates mission goals into a set of onboard activities to be performed and the dependencies among them; the Smart Executive, which constructs an explicit timeline of activities, and initiates and monitors the execution of those activities, and the Mode Identification and Reconfiguration system which continuously assesses

overall spacecraft state, diagnoses faults, and has the authority to command the spacecraft to the desired state.

Tom Gormley of Lockheed-Martin discussed work on Integrated Vehicle Health Management in the context of the next-generation space shuttle concept. Citing known cases where such onboard health monitoring could have been the key to avoiding launch scrubs or in-flight failures, Gormley emphasized the importance of systems analysis software engineering and described specific near-term targets for intelligent sensors on launch vehicles and orbiters. Examples include strain sensors on wing surfaces, temperature sensors, and acoustic sensors for detecting early damage.

Andy Mishkin of JPL reported on the autonomous capabilities of the Sojourner rover of the Mars Pathfinder mission. Sojourner receives commands from Earth once each Martian day and performs traverses, science target acquisitions and instrument placements autonomously. Obstacle detection and avoidance is accomplished with a laser-based system which illuminates nearby objects and infers their shape and proximity. Downlinked stereo images are utilized each day to plan way points for the next day's traverse. At the time of this writing, Sojourner's mission on Mars has been successfully completed.

Alan Schultz of the Naval Research Laboratory spoke to the core issue of how to test autonomy software, which is of a different nature and complexity from conventional onboard software and will likely require new software validation concepts. The key idea in this work is to employ genetic algorithms to explore the space of possible test scenarios, guided by human knowledge of fault classes as a starting point, but avoiding subtle biases which can result in inadequate coverage when humans generate the suite of test scenarios. The method has been evaluated for simulated autonomous landing of an F-14 on an aircraft carrier and was able to identify faults not anticipated by the designer.

David Kortencamp of NASA Johnson Space Center described an intriguing application of autonomy to robotic cameras performing inspection tasks of the Space Shuttle or Space Station or in support of astronauts performing Extra-Vehicular Activities (EVAs). This autonomy concept is based on a three-tiered architecture whose levels include a Skill Manager for low-level resource management and communication, a Sequencer which schedules and monitors specific activities and a Planner which determines the set of activities to achieve specific goals. Several experiments are planned on the Space Shuttle to validate this technology.

Maj. Richard Walker of DARPA offered a briefing on the Darkstar Unmanned Air Vehicle (UAV). The mission of this autonomous aircraft is reconnaissance in battlefield situations, with its purpose more succinctly

stated as: the right image to the right user at the right time at the right rate. The vehicle is being developed by Boeing and Lockheed-Martin. It achieves and maintains station autonomously and is commandable by line-of-sight or satellite communication for a range of wide area to spot coverage. Tests of the vehicle are underway.

Larry Matthies of JPL reported on work towards developing unmanned ground vehicles (UGVs) for the battlefield. The use of such assets, under a varying mix of remote control and autonomy, is a developing trend for battlefield technology into the next century. UGVs are being developed for capabilities such as reconnaissance, surveillance, target acquisition, obstacle breaching, countermine operations, and resupply. Tactical maneuvering and coordination among several UGVs are forms of autonomy that would be required for these vehicles to be truly effective.

Autonomy Technology for 2001

The third session of the workshop examined autonomous systems research and development for the near term, going out about five years.

Leon Alkalai of JPL described work at the Center for Integrated Space Microsystems, which is developing next-generation avionics for spacecraft. Targets include avionics on a chip, ultra-low power electronics, and evolvable hardware. The effort will also examine alternative computing technologies, such as DNA-based and quantum computing. CISM will work closely with the X2000 technology program at JPL, which is developing a generalized spacecraft architecture to include autonomy software.

Capt. Dan King of the USAF Phillips Laboratory gave an overview of the Improved Space Computer Program which is targeting advanced flight processors for military and commercial satellites. The effort at Phillips to develop spacecraft autonomy capabilities for reduced ground operations costs is interfaced with this program, as well as with autonomy development efforts within DoD and NASA.

Clark Chapman of the Southwest Research Institute spoke to applications of autonomy for planetary science. Autonomy has a role to play in those situations involving transient phenomena requiring timely decisions, those involving interactive operations in a remote location and those constrained by limited data rates. Onboard autonomy is not appropriate for the highest level cognitive functions of the scientist, but can support data acquisition, data reduction and classification of results in well-defined applications, thereby providing enhanced opportunities where scientists cannot possibly be involved otherwise.

Mark Hanson of TRW described an autonomy architecture which builds on prior work in next generation flight hardware and software architectures. The architecture includes a Mission Manager, which performs high-level planning, a System Manager, which performs fault management, and a Flight Manager which handles traditional spacecraft subsystem functions. The resulting system is commandable from at a high level and can both isolate and recover from unanticipated faults. The architecture has been evaluated on a number of simulated satellite scenarios.

Ethan Scarl of Boeing spoke to the status of automation and autonomy for the Space Station. Historically, in the early stages of Space Station development, there was a large effort to investigate automation and to develop some early prototypes. At this point, there are a number of control loops with automatic responses built into the current avionics system design. The overall architecture is conservative, however, with few responsibilities delegated to advanced automation. It is expected that automation and autonomy will be introduced on the Space Station, but primarily in peripheral systems like an external mobile inspection platform.

Glenn Yushimoto of Lockheed-Martin discussed a framework for intelligent data dissemination which begins with the ability to recognize and capture opportunistic science on space platforms using a suite of detection capabilities. Collected data is analyzed in several stages, for transmission from the space platform, reception and forwarding on the ground, and archiving to appropriate databases. Software agents evaluate and coordinate data traffic on the Web, leading to rapid publishing and efficient shunting of data to the useful destinations. The same capabilities could be used effectively in battlefield data management scenarios.

John Grefenstette of the Naval Research Laboratory reported on learning and adaptation in multi-agent systems. This has been a focused effort drawing on a range of learning techniques, from control theory to supervised learning to reinforcement learning. The research has also looked at issues in cooperating heterogeneous agents, where each agent may have different goals, and different problem-solving and learning capabilities. Agents should co-evolve to useful cooperative behaviors efficiently and robustly. Applications of interest to the Navy include surveillance, mine clearing and undersea equipment maintenance.

Ron Arkin of Georgia Tech gave an intriguing talk which examined robot design concepts from a suite of unusual viewpoints. Examples included: Imaginative robots, which simulate and explore the consequences of action before actually performing the action. Emotional robots, whose experience of frustration helps them triggers useful mode changes. Robots with hormones, which mediate internal communication and control functions. Robots which acquire skills using a form of learning analogous to immune system function.

And finally, in a rather startling example, a hybrid cockroach/robot system using a grafted microcontroller with potential for applications like pipe inspection.

Long-term Challenges and Benefits of Autonomy

The fourth and last session of the workshop presented bold, unfettered visions of where autonomy technology could reach, and what some of its ultimate payoffs might be.

Torrence Johnson of JPL, the Galileo project scientist, offered his views on autonomy from a scientist's perspective. He characterized science as conducted from space platforms as either "weather station" science or "telescope" science. Weather station science is mostly about reconnaissance, e.g., orbiter missions, and discoveries are made typically from the collected data. Telescope science is more targeted, with specialized instruments and objectives. Discoveries are made close to the data. The emerging class of *in situ* missions are good examples of this type. Johnson in general supports autonomy targeted for engineering functions of the spacecraft (e.g., navigation, fault protection) but cautioned that autonomy applied to science data processing or analysis should be used carefully, and is more appropriate for weather station science.

Richard Doyle of JPL presented a vision for the development of autonomy technology in which autonomy for science ultimately offers greater strategic value than autonomy for engineering or spacecraft functions, because science autonomy more directly enables new missions. Scientist-directed onboard software keeps the investigators in intimate contact with the spacecraft, allowing mission priorities to be evolved as scientific understanding of the remote environment evolves, through a combination of conventional algorithms, recognizers trained via machine learning techniques and knowledge discovery techniques, installed at launch time and uploaded during the mission. A number of ongoing scientist-defined projects in science autonomy were described, including natural satellite search and change detection on planetary surfaces.

Brian Williams of NASA Ames Research Center offered a vision for the development of future spacecraft and missions using a model-centric approach. The concept starts from the notion of model-based programming, where models not only capture knowledge, but also are composed directly to realize desired behaviors in the space system. A model-based autonomy kernel can be realized, combining reactive and deductive capabilities, which supports such useful behaviors as anticipation, self-modeling, adaptation, information seeking, and collaboration. Common modeling tools will be

essential in realizing this vision, as will new validation techniques, which may themselves draw on model-based concepts.

Robert Ferraro of JPL described the NASA Remote Exploration and Experimentation Program. The goal of this program is to move scaleable supercomputing technology into space, driving the development of low-power, fault tolerant, scaleable computing technologies in partnership with industry and the science community. Advances in this technology area will support concurrent advances in autonomy technology and high data rate sensors and instruments, playing an important part in enabling future science missions which will be characterized by *in situ* investigations.

Gregg Swietek of Lockheed-Martin spoke to the need to view technologies like autonomy across the entire mission lifecycle from initial design through to deployment and operations. He emphasized the value of simulation-based design in particular for maximizing cost benefit across a single lifecycle and amortizing costs across multiple uses of a technology. Autonomy can contribute during initial design by enabling new mission concepts. The use of autonomy within a mission must be explicit when system trades are conducted during early detailed design, for much of the benefit of autonomy is at the system level.

David Collins of JPL reported on activities of the Microspacecraft Systems Technology effort. Projects include applications of machine vision to dim star recognition for navigation purposes, an approach to sample selection on the surface of Mars based on fuzzy logic and evolvable hardware, neural networks for spin vector identification in spacecraft control, neural networks for target recognition and tracking, optical processing for precision landing on planetary surfaces, and genetic algorithms for flight path optimization.

Gerald Sussman of MIT, in a talk entitled "The Future As I See It," presented a perspective on technology evolution as the development of different kinds of prosthetics, where prosthetic is taken in its general sense as a compensator or amplifier for an ability which has been compromised or is inadequate for the task at hand. Different eras have different views on what form of machine assistance is most useful. The industrial revolution might be taken as the successful development of mechanical prosthetics. The medical prosthetics emerging today may be just the vanguard of a more general class of biological prosthetics. Sussman explored bold concepts for intelligence prosthetics of the future, the logical successor in the sequence of prosthetics development.

Banquet Speakers

The workshop was graced with two exceptional banquet speakers, both of whom represented the theme of celebrating 1997 as the fictional birthyear of the HAL 9000 computer. The first speaker was none other than the creator of HAL 9000, the author of 2001: A Space Odyssey, Dr. Arthur C. Clarke. Dr. Clarke prepared a videotaped greeting to the workshop attendees from his residence in Sri Lanka. The second speaker was Prof. David Stork of Stanford University. Prof. Stork is the author of HAL's Legacy, which examines the technology prophecies of 2001 from the perspective of the present, providing a number of delightful surprises.

Arthur C. Clarke organized his address to the workshop attendees around a set of reminiscences about HAL 9000 and Clarke's own personal interactions with NASA. He recounted how his career has spanned the origins of spaceflight, from the development of rocketry theory, through the realization of spaceborne telecommunications satellites, a concept he first articulated, to the active exploration of the solar system. He lamented the common misinterpretation of the basis of the name HAL, reminding us that the correct derivation is Heuristically Programmed ALgorithmic Computer, one which "has the best of both worlds." He bemused us with how casual, tongue-in-cheek remarks (in this case, regarding the image of a "face" on Mars) are quickly taken up by the less responsible media. He challenged us to explore Jupiter's moon Europa, speculations concerning which have appeared in his recent fiction. Clarke left us with the following paraphrase of Descartes, perhaps to be uttered someday by an intelligent machine: "I think, therefore I am, I think."

David Stork gave a wonderfully engaging talk on the topic of the remarkable prescience regarding computer science-based technologies to be found in both the novel and the movie 2001: A Space Odyssey. His main motivation for writing HAL's Legacy was aesthetics. He wanted to use his knowledge of science and technology to offer a deeper understanding of the film and book, thereby enabling the reader to see them more sensitively, with a more educated eye. Stork reexamined Clarke's and Kubrick's vision and found that in some cases, technology maturity has fallen short of their predictions, but in others, has surged notably ahead. Computer graphics, supercomputing hardware and computer reliability have surpassed the vision of 2001. On the other hand, speech recognition, language understanding, common sense, and planning have all fallen short -- far short. Clearly, the general level of artificial intelligence exhibited by HAL 9000 is not likely to be achieved by the year 2001. On the other hand, a computer has now defeated the world chess champion. The conversational interactions between HAL and the human characters Poole and Bowman are loaded with implications about speech generation, speech analysis, and facial expression analysis technology. Stork unearthed from the archives of AT&T what must be the inspiration for some

of HAL's discourse: a tape from the 1960s of an early speech generator reciting the verse of the song "Daisy." He also described current work at the MIT Media Lab and elsewhere on inferring the emotional state of a speaker (as an input to semantic analysis of speech) from inflection analysis of the speech signal and visual analysis of facial images. Stork's book contains many additional fascinating examples and insightful analyses.

University Design Competition on Aerobots

Aerobots are a new space platform concept which combines some of the best aspects of orbiter-style missions and surface-style missions. Specifically, an aerobot is designed to exploit the diurnal thermal cycle of a planetary environment by going aloft once a sol (a sol is the term assigned to the day cycle in the local planetary environment) and landing once a sol. In this way, an aerobot achieves in part the wide coverage aspects of an orbiter mission, which can survey an entire planetary surface, along with the *in situ* exploration aspects of a surface mission, such as those executed by lander and rover combinations, where scientific experiments are conducted in direct interaction with the planetary environment. Although it is possible to predict with some accuracy where an aerobot may land next (with knowledge of prevailing planetary wind patterns for example) aerobots sample the planetary environment in a stochastic manner, making it nearly impossible to return to a site after leaving. Aerobots are being conceived for exploration wherever planetary atmospheres are present, including Venus, Mars, Jupiter, and Saturn's moon Titan.

Aerobots will require a significant degree of autonomy. Communication will be very problematic within the atmospheres at destinations like Venus and Jupiter. If successful scientific missions are to be achieved there, the aerobot platform must be able to grapple with uncertainty again and again and continue to plan and execute the mission while going long periods. Path planning with a significant stochastic element will be only one of several unique challenges.

The design competition on aerobots was announced at this workshop, targeted to the university community. The intent is to start a cycle where the submissions from the previous design competition are reviewed at each Highly Autonomous Systems Workshop, and a new design competition is announced. Reid Simmons of Carnegie-Mellon University is acting as the coordinator for the aerobot design competition.

The Future of Autonomy

In the story 2001: A Space Odyssey several decades ago, whether through brilliant foresight or the whimsy of time, Arthur C. Clarke correctly predicted the turning of the millennium as a pivotal moment in the development of highly autonomous systems. He also predicted the momentous impact this would have on our future -- a development that would change forever our views of exploration and the bounds of our experience. In the Highly Autonomous Systems Workshop we gathered, not only to celebrate this great act of prescience, but also to share our collective experiences and vision for autonomy.

This workshop demonstrated in one presentation after another the broad interest and investment in this technology present today throughout the aerospace, defense, scientific, and exploration communities. It demonstrated that the ideas, computational power, and conviction to make it work are in place. And it demonstrated that advanced autonomy is viewed seriously as a *practical* answer to real and pressing needs. This unprecedented confluence of need and readiness heralds an era of enormous possibilities.

Highly autonomous systems will greatly extend the safe and efficient exploration of space by enabling probes to hostile and unpredictable places. They will help us understand our own fragile planet from ocean floor to volcanic peak by guiding fleets of explorers and scrutinizing inexhaustible sources of data. They will enhance our national defenses by placing only artificial eyes and ears in harm's way. They will help save lives in space, in the air, and most importantly, on the ground by providing warnings of danger for everything from malfunctioning systems to tsunamis. All of these needs are compelling. Our success in addressing them is not of mere academic interest, but rather serves a vital societal role.

True success, therefore, must be measured in the eagerness of the world to adopt autonomy. Yet ironically, the greatest obstacle to this progress is autonomy's own basic nature. The long term vision of 2001 cast intelligent machines, not merely as tools, but more significantly, as partners to the human endeavor, capable of deliberate independent action. This is the essence of the word "autonomy", but it is what skeptics fear most. Independent action is taken as action that is out of control. Moreover, it is often viewed as a usurpment of human volition -- an expensive way to do the wrong thing.

The ultimate challenge to highly autonomous systems will therefore be the happy union of control and independence we are able to concoct, in order that this technology should find an open invitation to wide use. To this end future workshops will continue to concentrate on this imperative but visionary aspect of autonomous systems and their successful injection into

real world practical applications. We will follow developments from concept to realization in the field to hard lessons learned, and we will chart the purposeful advancement of the technology, providing a forum for objective appraisal.

The future of autonomy is in your hands. We look forward to hearing from all of you at our next workshop.